ABSTRACT

Prior to searching a multidimensional feature space populated with data objects, each dimension in the feature space is divided into a number of intervals. When a query is received, a single interval that is overlapped by the query is selected from each dimension. A reduced set of data objects is then selected that includes only those data objects that overlap the selected intervals. This reduced set of data objects, rather than the entire set of data objects in the feature space, is then used to determine matches for the query.